Simulation-Based Synthesis of Personalised Therapies for Colorectal Cancer

Marco Esposito, Leonardo Picchiami

Computer Science Dept., Sapienza University of Rome, Italy

Abstract

In this short paper we present preliminary results on computing, *in silico*, personalised therapies for Colorectal Cancer (CRC), one of the deadliest form of tumour for adult humans. We exploit a recent SBML (Systems Biology Markup Language) model of the tumour growth, which also models the PharmacoKinetics/Pharmacodynamics (PK/PD) of two immunotherapic drugs that may be used in combination treatments.

Keywords

In Silico Clinical Trials, Precision Medicine, Personalised Therapies, Automated Synthesis

1. Introduction

The recent availability of quantitative models of the human patho-physiology (Virtual Physiological Human, VPH) has inspired and facilitated new approaches to the design of pharmacological treatments and the safety and efficacy assessment of therapies and biomedical devices. These kinds of approaches, collectively referred to as In Silico Clinical Trials (ISCTs), hold the promise to enable *precision medicine*, in which Artificial Intelligence (AI) methods support the design of *personalised* therapies and the assessment of their efficacy by means of simulation (*i.e.*, *in silico*). ISCTs require computational VPH models that take into consideration the physiological differences between different human individuals and that capture the kinetics and dynamics of pharmacological drugs. Quantitative VPH models are typically defined as *hybrid systems*, where the dynamics is described by systems of Ordinary Differential Equations (ODEs) and the inter-patient variability is encoded by *parameters*. These models are often designed and distributed in open-standard modelling languages, often the Systems Biology Markup Language (SBML).

2. Modelling

The automatic synthesis of therapies that may involve several drugs is a computationally complex task. In fact, it is necessary to determine which drugs to use and their associated dosing regimen, *i.e.,* the drug amounts and the frequency of administration. Also, as it is the case

for most pathologies and drugs, not all possible therapies are actually feasible; most notably, there exist known constraints that limit the quantity of drugs that may be administered in a certain period, in order to deal with drug toxicity. In the general case, the number of possible therapies is infinite. We study an approach to the problem that is based on the parametrisation of therapies. In order for the approach to be both feasible and effective, the number of parameters must be small enough to enable an efficient search over the possible therapies but still large enough to allow the modelling of all (or most of) the therapies of interest. Our goal, then, is the synthesis of therapies (assignment to the parameters) that meet a given set of constraints and optimise a user-defined objective function, which usually combines a set of performance metrics that include *efficacy* (*i.e.*, how well the therapies cures the pathological condition) and the total quantity of administered drugs (generally to be minimised, so to reduce costs and health risks). We propose a method that uses VPH models described in the SBML language and that defines the whole optimisation problem inside the model itself. In particular, the parameters of the treatment will be modelled as additional parameters of the VPH model and the objective function and the therapy constraints will be modelled as observable model outputs. The fact that the whole problem is defined in pure SBML, a standard language supported by a large suite of software and with a large community, has the goal of showing that computational methods have the potential to be used in an easy and immediate way in clinical contexts.

3. Computing personalised therapies for Colorectal Cancer

This section describes the steps we followed to set up an ISCT to evaluate the effectiveness of our approach in the synthesis of personalised therapies for the CRC model proposed in [1].
Generation of a virtual population. The starting point to perform an ISCT is the availability of a complete and representative population of virtual patients (VPs), *i.e.,* assignments to the VPH model parameters. The CRC model presents 23 real-valued parameters that encode the inter-patient variability. We used the approach described in [2, 3, 4] to generate a population of VPs that are of interest (*i.e.*, do actually develop a tumour) and show evolution of the model observables that are *physiologically admissible* (*i.e.*, do not violate the laws of biology).
Therapy modelling. We modelled 60-week-long therapies as assignments to a set of 32 parameters. For each of the two drugs, *atezolizumab* and *cibisatamab*, each one of 15 parameters governs the amount of drug that can be administered during a 28-day period and one additional parameter defines the number of weeks between two consecutive administrations of the drug.
Therapy constraints and objective function. We modelled two constraints, each of which limits the total amount of each drug cumulatively administered in each 28-day period according to known toxicity levels and past clinical trials (https://clinicaltrials.gov/ct2/show/NCT03866239). The objective function combines the total amount of administered drug with a measure of *inefficacy* of the treatment. This last measure is computed based on the volume of the tumour at the end of the treatment and its initial volume. An high value for the inefficacy measure means that the therapy is not able to keep the tumour growth under control.

4. Experimental results

We chose to carry out the experimental evaluation of our approach using COPASI [5], one of the most well-known software tools supporting both plain simulation and simulation-based parameter optimisation of SBML models via iterative improving algorithms. The starting point of our ISCT was the random sampling of 35 VPs from the previously computed complete population. We compared the optimisation performance of 3 algorithms implemented in COPASI: the standard Genetic Algorithm (GA), Particle Swarm Optimisation (PSO), and the Levenberg-Marquardt (LM) algorithm, a gradient-descent based method that combines Steepest Descent and the Newton Method. The hyper-parameters of the algorithms were chosen as to perform the optimisation in a reasonable time for every patient (within 2 hours on an Intel Xeon E5430 @ 2.66GHz (8 cores) 32GB RAM machine). In order to compare the quality of the personalised therapies synthesised by each algorithm with a common baseline, we considered the dosing regimen of [1] as a a a reference. Given a therapy, we define the Drug Amount Percentage as $\frac{t_{\text{amount}}}{r_{\text{amount}}} \times 100\%$, where t_{amount} and r_{amount} are the total amounts of drugs administered by the given therapy and the reference one, respectively. For each VP, the Inefficacy Percentage for a given therapy is defined as $\frac{t_{\text{ineff}}}{r_{\text{ineff}}} \times 100\%$, where t_{ineff} and r_{ineff} are the inefficacy values for the given and the reference therapy, respectively. The objective function is a linear combination of these two metrics, where the coefficients are chosen so to balance the search for effective treatments and the minimisation of administered drug amounts. In our experiments, only 11 VPs out of 35 showed a response to the drugs. This is in agreement with the results from [1].

For such 11 patients, LM was not able to find a therapy that improves the reference one, while GA and PSO show, on average, good results. The average reduction of the amount of administered drug is *as high as* 96.9% for GA and 98.62% for PSO, while the average reduction of the inefficacy is around 35% for both algorithms. Nonetheless, such personalised therapies still manage to reduce the tumour growth significantly.

5. Related Work

Many attempts have been proposed in the literature to solve *large* optimisation problems defined via logic-, automata- or constraint-based formalisms (*e.g.*, [6, 7, 8, 9, 10, 11, 12] among others). However, such approaches cannot be applied when the problem model (a complex ODE-based VPH model, as in our case) cannot be accurately defined within such formalisms and is available only as a simulator. Indeed, although such VPH models are hybrid systems whose inputs are discrete event sequences [13, 14], to find an optimal treatment means to find an optimal plan in hybrid domains. Although symbolic approaches exist to model and solve planning problems in hybrid domains [15, 16], the typical complexity of the ODEs of clinically-relevant VPH models makes such models out of reach for them, and appoints numerical integration as the only viable means to compute (black-box) the model evolutions under a given input function.

The synthesis of personalised therapies exploiting black-box VPH models is addressed in, *e.g.*, [17, 18]. In [19, 20] the authors propose a intelligent backtracking simulation-based algorithm guided by multiple heuristics to seek, on a patient *digital twin*, an optimal robust personalised treatment for assisted reproduction, an area with many factors hard to control [21, 22, 23] and

for which treatment personalisation is crucial for success.

One of the main problems in system biology is the estimation of unknown model parameters that fit a series of experimental data. Various optimisation algorithms are studied in [24, 25, 26] and applied in real-world case studies in, *e.g.*, [27, 28, 29]. Many of the available tools rely on SBML simulators (see www.sbml.org). Among such simulators is SBML2Modelica [30], which focuses on the interoperability between system biology and (hybrid) CPS domain by translating SBML models to Modelica and the FMI/FMU open standard. This enables the seamless exploitation of tools and methodologies already established for CPS optimisation and verification, in particular backtracking-based search and optimisation via the efficient storing and retrieval of intermediate simulator states [31], verification of closed-loop systems also in presence of uncontrollable events (*e.g.*, [32, 33, 34, 35, 36, 37, 38]).

6. Conclusions

The good results of the GA and PSO algorithms show the potential of the approach in the synthesis of personalised therapies. We interpret the failure of the LM algorithm as a clue of the fact that purely gradient-based optimisation is not suited for this problem, due to the strong constraints enforced on the therapies. Conversely, population-based algorithms show good results thanks to their ability to widen their focus throughout the search space.

References

[1] H. Ma, *et al.*, Combination therapy with t cell engager and pd-l1 blockade enhances the antitumor potency of t cells as predicted by a qsp model, J Immun Cancer 8 (2020).

[2] E. Tronci, *et al.*, Patient-specific models from inter-patient biological models and clinical records, in: FMCAD 2014, IEEE, 2014.

[3] T. Mancini, *et al.*, Computing biological model parameters by parallel statistical model checking, in: IWBBIO 2015, *LNCS* 9044, Springer, 2015.

[4] S. Sinisi, *et al.*, Complete populations of virtual patients for in silico clinical trials, Bioinf 36 (2020).

[5] S. Hoops, *et al.*, Copasi—a complex pathway simulator, Bioinf 22 (2006).

[6] M. Cadoli, *et al.*, Combining relational algebra, SQL, constraint modelling, and local search, Theor Pract Logic Progr 7 (2007).

[7] M. Cadoli, *et al.* SAT as an effective solving technology for constraint problems, in: ISMIS 2006, *LNCS* 4203, Springer, 2006.

[8] T. Mancini, *et al.*, Combinatorial problem solving over relational databases: View synthesis through constraint-based local search, in: SAC 2012, ACM, 2012.

[9] T. Mancini, *et al.*, An efficient algorithm for network vulnerability analysis under malicious attacks, in: ISMIS 2018, Springer, 2018.

[10] T. Mancini, *et al.*, Optimal fault-tolerant placement of relay nodes in a mission critical wireless network, in: RCRA 2018, *CEUR* 2271, 2018.

[11] Q. Chen, *et al.*, MILP, pseudo-boolean, and OMT solvers for optimal fault-tolerant placements of relay nodes in mission critical wireless networks, Fund Inf 174(3-4) (2020).

[12] I. Melatti, *et al.*, A two-layer near-optimal strategy for substation constraint management via home batteries, IEEE Trans Ind Elect (2021).

[13] T. Mancini, *et al.*, Anytime system level verification via random exhaustive hardware in the loop simulation, in: DSD 2014, IEEE, 2014.

[14] T. Mancini, *et al.*, SyLVaaS: System level formal verification as a service, in: PDP 2015.

[15] M. Fox, *et al.*, Modelling mixed discrete-continuous domains for planning., JAIR 27 (2006).

[16] S. Bogomolov, *et al.*, PDDL+ planning with hybrid automata: Foundations of translating must behavior., in: ICAPS 2015, AAAI, 2015.

[17] T. Cassidy, *et al.*, Determinants of combination gm-csf immunotherapy and oncolytic virotherapy success identified through in silico treatment personalization, PLoS Comp Biol 15 (2019).

[18] A. L. Jenner, *et al.*, Optimising hydrogel release profiles for viro-immunotherapy using oncolytic adenovirus expressing il-12 and gm-csf with immature dendritic cells, Appl Sci 10 (2020).

[19] T. Mancini, *et al.*, Computing personalised treatments through in silico clinical trials. A case study on downregulation in assisted reproduction, in: RCRA 2018, *CEUR* 2271, 2018.

[20] S. Sinisi, *et al.*, Optimal personalised treatment computation through in silico clinical trials on patient digital twins, Fund Inf 174(3−4) (2020).

[21] B. Leeners, *et al.*, Lack of associations between female hormone levels and visuospatial working memory, divided attention and cognitive bias across two consecutive menstrual cycles, Front Behav Neur 11 (2017).

[22] M. Hengartner, *et al.*, Negative affect is unrelated to fluctuations in hormone levels across the menstrual cycle: Evidence from a multisite observational study across two successive cycles, J Psychol Res 99 (2017).

[23] B. Leeners, *et al.* Associations between natural physiological and supraphysiological estradiol levels and stress perception, Front Psycol 10 (2019).

[24] L. Schmiester, *et al.*, Efficient gradient-based parameter estimation for dynamic models using qualitative data, Bioinf (2021).

[25] A. F. Villaverde, *et al.*, Benchmarking optimization methods for parameter estimation in large kinetic models, Bioinf 35 (2019).

[26] A. Yazdani, *et al.*, Systems biology informed deep learning for inferring parameters and hidden dynamics, PLoS Comp Biol 16 (2020).

[27] O. D. Sánchez, *et al.*, Parameter estimation of a meal glucose−insulin model for tidm patients from therapy historical data, IET Sys Biol 13 (2019).

[28] T. Chen, *et al.*, Optimal dosing of cancer chemotherapy using model predictive control and moving horizon state/parameter estimation, Comp Meth Progr Biomed 108 (2012).

[29] M. Raissi, *et al.*, On parameter estimation approaches for predicting disease transmission through optimization, deep learning and statistical inference methods, Lett Biom 6 (2019).

[30] F. Maggioli, *et al.*, SBML2Modelica: Integrating biochemical models within open-standard simulation ecosystems, Bioinf 36 (2020).

[31] S. Sinisi, *et al.*, Reconciling interoperability with efficient verification and validation within open source simulation environments, Simul Model Pract Theory 109 (2021).

[32] T. Mancini, *et al.*, SyLVaaS: System level formal verification as a service, Fund Inf 149(1−2) (2016).

[33] T. Mancini, *et al.*, Anytime system level verification via parallel random exhaustive hardware in the loop simulation, Micropr Microsys 41 (2016).

[34] T. Mancini, *et al.*, On minimising the maximum expected verification time, Inf Proc Lett 122 (2017).

[35] T. Mancini, *et al.*, User flexibility aware price policy synthesis for smart grids, in: DSD 2015.

[36] T. Mancini, *et al.*, Demand-aware price policy synthesis and verification services for smart grids, in: SmartGridComm 2014, IEEE, 2014.

[37] T. Mancini, *et al.*, Parallel statistical model checking for safety verification in smart grids, in: SmartGridComm 2018, IEEE, 2018.

[38] T. Mancini, *et al.*, Any-horizon uniform random sampling and enumeration of constrained scenarios for simulation-based formal verification, IEEE TSE (2021). To appear.